

GALAXY PACKAGE FOR N -BODY SIMULATION

J. A. SELLWOOD

Dept. of Physics & Astronomy, Rutgers University,
136 Frelinghuysen Road, Piscataway, NJ 08854, USA

ABSTRACT

This posting announces public availability of the **GALAXY** software package developed by the author over the past 40 years. It is a highly efficient code for the evolution of (almost) isolated, collisionless stellar systems, both disk-like and ellipsoidal. In addition to the N -body code **galaxy**, which offers eleven different methods to compute the gravitational accelerations, the package also includes sophisticated set-up and analysis software. This paper gives an outline of the contents of the package and provides links to the source code and a comprehensive on-line manual. While not as versatile as tree codes, the particle-mesh methods in this package are shown, for certain restricted applications, to be between 50 and 200 times faster than a widely-used tree code.

1. INTRODUCTION

This publication announces the **GALAXY** software package¹ written almost exclusively by the author for collisionless simulations of galaxies. I have developed the code over a period of almost 4 decades, and continue to make refinements; it has now reached version 14. Pieces of the code have been described in a number of papers in the literature over the years (James & Sellwood 1978; Sellwood 1981, 1985; Sellwood & Athanassoula 1986; Earn & Sellwood 1995; Sellwood & Valluri 1997; Debattista & Sellwood 2000; Sellwood 2003, *etc.*) The full source code is available, and may be downloaded and installed as described in §1.4. The code may be copied freely (see §1.2), but people publishing results that use the software it provides should acknowledge the source and cite this paper.

An on-line manual for use of the code is maintained at <http://www.physics.rutgers.edu/~sellwood/manual.pdf>. Since this is comprehensive, I give only an outline here.

The quality of any collisionless N -body simulation is almost entirely limited by the number of particles employed, which places a high premium on efficiency. For many applications of interest, the 3D methods described here are tens to hundreds of times faster than the popular tree codes, and the 2D methods have a still greater performance advantage. However, speed comes at the cost of reduced flexibility and increased danger of compromised results from inappropriate use. While the code is neither as versatile nor as easy to use as other publicly available packages, the careful user should find the increased quality of the results obtainable is ample compensation for the effort required to learn how to use it.

1.1. Contents of the package

To make use of an N -body code, one must first set up the model to be evolved and afterwards analyze the behavior of the simulation in order to draw conclusions. The **GALAXY** package therefore supplies three integrated pieces of code:

- A The most obvious, and in many ways the easiest part to develop and describe, is the code to compute the evolution of an N -body system. The program **galaxy** advances a pre-created file of particles in time for as long as is desired. This is the core of the code, where efficiency is key. At the end, and at intermediate times if desired, it creates a file with the full information about the state of the system, from which the evolution can be continued.
- B Analysis of the results is both more difficult to develop and to describe, largely because almost every project has a different scientific purpose, and therefore requires some piece of analysis that could be unique. However, most projects also make use of standard analyses of the system state and the code provides many such options. Unlike most other publicly available codes, these standard analyses in the **GALAXY** package are computed “on-the-fly”, and the results at each analysis step are saved in a file for later examination. The huge advantage of analysis on-the-fly is that the results are much more compact and can therefore be saved at frequent intervals without the inconvenience of large data storage and time-consuming analysis programs. The principal disadvantage is, of course, that a simulation must be rerun if a quantity was not measured in the first run. However, the analysis options are now well enough developed that this happens but rarely.

¹ © 2014. The copyright for this package is owned by the author

C By far the most difficult part of a successful N -body project is the creation of the desired initial set of particles from which to start the simulation. Assigning initial positions and velocities of a specified system that is to be in equilibrium is generally hard, and it becomes especially difficult for a system, such as a disk galaxy, with multiple mass components: a disk, bulge, and a halo. The code package includes sophisticated software that can be used in some cases, but naturally not every possible option is anticipated.

1.2. License

GALAXY is a package of free software, distributed under the GNU General Public License.² This implies that the software may be freely copied and distributed. It may also be modified as desired, and the modified versions distributed as long as any changes made to the original code are indicated prominently and the original copyright and no-warranty notices are left intact. Please read the General Public License² for more details. Note that the author retains the copyright on the code.

1.3. Source code

The code is written in standard Fortran (f90) and should compile without difficulty. The basic **galaxy** program uses FFTs from the open access package **FFTPAK** (Schwarztrauber 1982) that are bundled with the source code. The post-run analysis software requires the open access package **PGPLOT**.³

However, some parts of the code do make use of the NAG library,⁴ which is a commercial package. Many institutes and universities may have it and even a very old version would be adequate. The package attempts to manage this aspect in as user-friendly manner as possible through the use of an environment variable.

The initialization program **begin** and the main N -body **galaxy** program should execute without the NAG library for most applications. **Thus most N -body simulations can be run without the NAG library.** Other goodies, such as parts of the analysis code, and essentially all the set-up software, make more extensive use of the NAG library and without it, those wishing to use these powerful tools will have to make a significant programming effort to substitute routines from no-cost libraries.

The main programs **begin**, **start**, **galaxy**, and **finish** do not make use of MPI calls and can therefore be used only for running on a single processor. There is a separate set of executables **begin_mpi**, **start_mpi**, **galaxy_mpi**, and **finish_mpi** that enable calculations on multiple processors, which therefore require a system with some version of MPI. Even though these executables must be linked with MPI, they can also be used for running on a single processor if desired. Throughout the documentation, the names of these programs are used without the extension **_mpi** to imply both versions.

1.4. Downloading and compiling the code

The installation assumes **PGPLOT**³ and a version of MPI are installed. The NAG⁴ library is optional but, without it, many powerful features will be unavailable.

Proceed with the following steps:

1. To download the software, go to the website

`http://www.physics.rutgers.edu/galaxy`

Enter your email address, click **download** and the file **galaxy.tgz** will download to your system automatically. (The email address will enable you to be notified when updates are available.)

2. This gzipped ~ 6 MB tar file can be moved from your Download directory to some convenient location in your directory tree and then unpacked with the command:

`tar xzf galaxy.tgz`

and all the files of the package will be created in the directory **pckg** and subdirectories thereof. The scripts and Makefile assume that the user will be using **gfortran**, and will need to be edited if a different compiler is desired.

² see <http://www.gnu.org/copyleft/gpl.html>

³ available from <http://www.astro.caltech.edu/~tjp/pgplot/>

⁴ see <http://www.nag.co.uk/numeric/fl/FLdescription.asp>

3. Build libraries from the source code. (These two `rebuild` scripts use `gfortran` by default.)

```
cd pkg/lib14
./rebuild
cd ../sfftpak
./rebuild
```

These scripts should create the two object libraries `main14.a` and `sfftpak.a` in the main directory `pkg`.

4. To run in parallel, the package will need an implementation of MPI. The scripts and Makefile assume `OpenMPI`⁵ and they will need to be edited if a different installation of MPI is to be used. The mpi compiler is assumed to be `/opt/openmpi/gnu/bin/mpif90`

```
cd ../lib14_mpi
./rebuild
```

This script should create the additional object library `main14_mpi.a` in the directory `pkg`.

```
cd ../progs
```

5. To inform the compiler of the location of the PGPLOT object library, edit the line that begins `pgplot_lib =` in the Makefile (in the `progs` subdirectory).
6. Decide whether the NAG library is available. Skip to step 7 if NAG is unavailable, otherwise there are two additional steps:
 - a Edit the line that begins `nag_lib =` in the Makefile (in the `progs` subdirectory) to inform the compiler of the location of the NAG object library.
 - b Set the environment variable `NAG` with the following command:

```
setenv NAG nag
```

If this environment variable is not set, or if it is set to any other value, then a restricted fraction of the code will be compiled. While these programs are generally expected to execute normally, they could terminate gracefully with an explanatory message if a NAG routine is required for some particular application.

7. Compile the code:

```
make
make clean
```

The `make clean` will not only tidy up the directory, but will move all the executables to the user's directory `$HOME/bin`, which is presumed to be in the user's `PATH`. To update your path, issue the command:

```
rehash
```

Depending upon whether NAG is available, the following executables of the `GALAXY` package should now be ready for use:

```
analys, begin, begin_mpi, cold, compress, corrp1t, dfiter, dflook, finish, finish_mpi, galaxy,
galaxy_mpi, isotropy, merge, modefit, plotpft, ptest, ptest_mpi, smooth, start, start_mpi, weed
```

Without NAG, only the following restricted set of programs can be used:

```
analys, begin, begin_mpi, finish, finish_mpi, galaxy, galaxy_mpi, merge, ptest, ptest_mpi
```

Executables without the `_mpi` suffix are for single processors only. The use of these programs is described in the on-line manual <http://www.physics.rutgers.edu/~sellwood/manual.pdf>.

1.5. Benchmark tests

Sellwood & Debattista (2006) reported that `galaxy` ran ~ 37 times faster than the non-public tree code `PKDGRAV` (Stadel 2001) on a single processor. Here I report further benchmark tests against the public tree code `Gadget-2` (Springel 2005). While `Gadget-2` has been used quite extensively to simulate isolated galaxies, it was designed for cosmic structure formation calculations that cannot be simulated with `galaxy`. This comparison is simply to highlight the relative efficiency of `galaxy` for isolated galaxies, and is in no way intended as a criticism of `Gadget-2`.

⁵ available from <http://www.open-mpi.org/software/ompi/v1.8/>

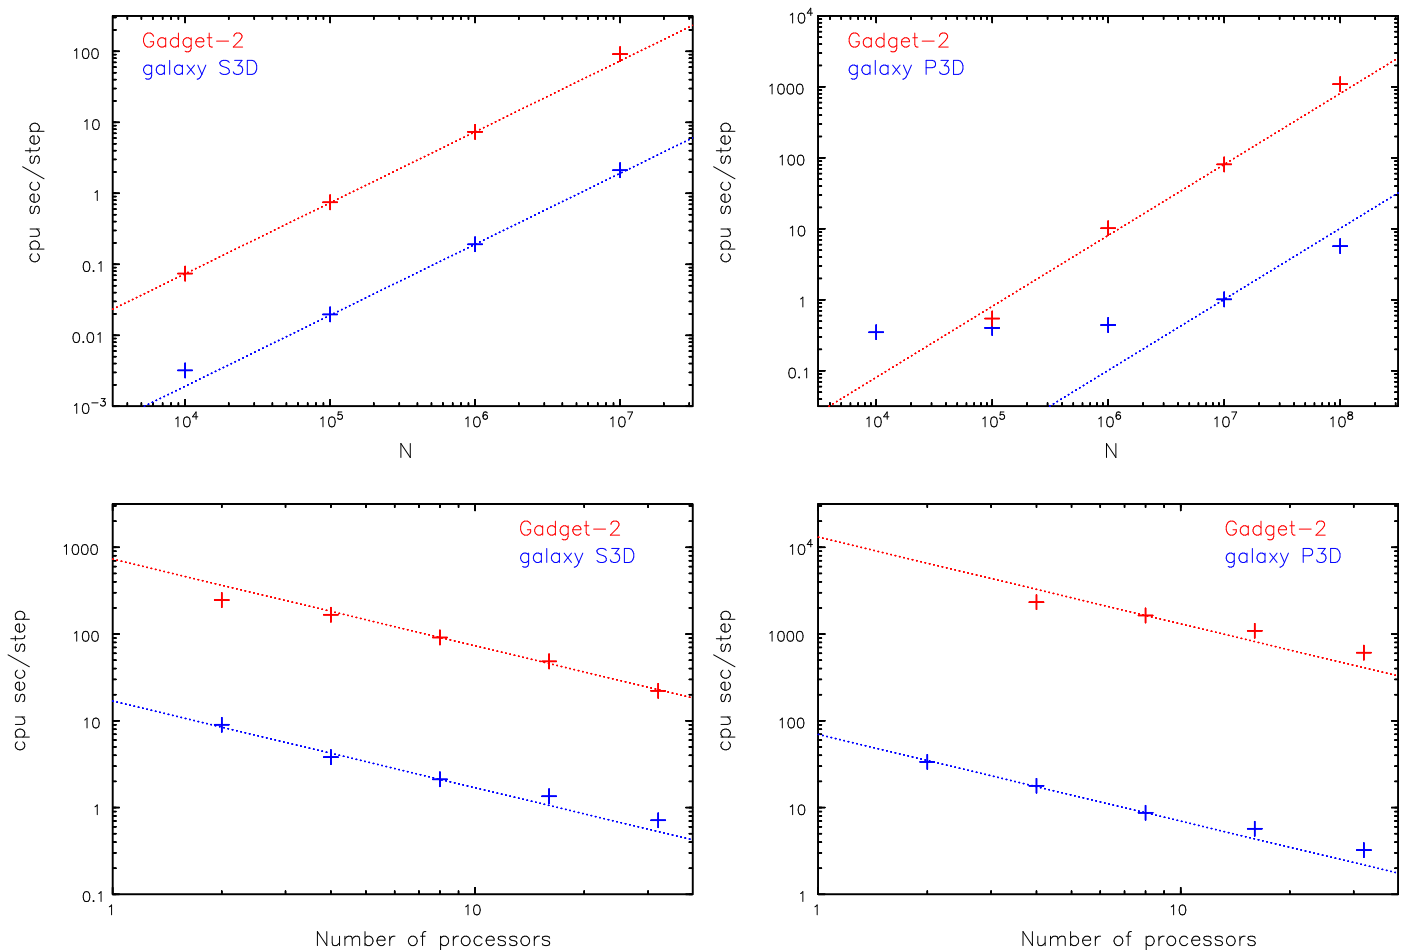


FIG. 1.— CPU timings for **galaxy** and **Gadget-2**. The left panels are for a spherical model, the right panels for a disk. The upper left panel shows the scaling with particle number N for 8 cores while the upper right uses 16 cores. The lower left panel shows scaling with the number of processors for $N = 10^7$, while the right is for $N = 10^8$. The dotted lines indicate slopes of ± 1 , and are not fits to the data.

Fig. 1 shows the measured cpu (wallclock) time for identical files of equal-mass star particles in simulations with **galaxy** and **Gadget-2**. The timings shown are cpu seconds per step averaged over 100 time steps. The upper panels show scaling with N for 8 cores (left) and 16 cores (right), while the lower panels show scaling with the number of cores for $N = 10^7$ (left) and $N = 10^8$ (right). All other numerical parameters were held fixed as the particle number and number of cores were varied. The opening angle used in **Gadget-2** was 0.5 radians and the particle softening length was $0.05a$.

The left-hand panels are for an equilibrium Plummer sphere, of core radius a , using the spherical grid option of **galaxy** with surface harmonics $0 \leq l \leq 4$ included in the force determination. There were 4 time-step zones, and the grid was recentered every 8 time steps. The right-hand panels are for a thickened KT disk run using a cylindrical polar grid of size $N_R \times N_a \times N_z = 86 \times 128 \times 125$, with sectoral harmonics $0 \leq m \leq 8$ included in the force determination. In this case, there were 3 time-step zones, and the grid was not recentered.

As may be seen from the Figure, **galaxy** runs ~ 50 times faster than **Gadget-2** for spherical models consistently for all N and regardless of the number of processors employed. For the disk model, the speed advantage of **galaxy** over **Gadget-2** gradually improves as N is increased. Because there is a fixed overhead for the grid force determination, the running time for **galaxy** barely increases until $N \gtrsim 10^6$, and for very large N , **galaxy** can be $\gtrsim 200$ times faster than **Gadget-2**.

1.6. Disclaimers

This free software comes with ABSOLUTELY NO WARRANTY. Users are welcome to make use of it, change it, and to redistribute it under certain conditions – see <http://www.gnu.org/licenses/> for details.

It is not expected that a would-be user could download the software provided and quickly run his/her desired simulation using it. The basic **galaxy** program should prove relatively easy to use as a “black box” and I provide a few test cases

whose results the user should be able to reproduce without additional coding. But many possible pitfalls could vitiate results from hasty use in a new application and I anticipate that a capable prospective user would need to gain many weeks of experience using the package before feeling sufficiently confident that results are publishable.

While the author welcomes reports of bugs in the distributed software, which he will attempt to fix, he is unable to provide help to modify it for some other application. The user should run the standard tests to ensure that the package has been properly installed, and report to the author if any of these tests fail on the user's system. But he cannot offer help beyond that point.

ACKNOWLEDGMENTS

The author wishes to thank all the many collaborators and graduate students with whom he has worked over the years, many of whom have made helpful suggestions to improve or add to the software. Especially noteworthy contributions have been made by my thesis advisor Richard James, whose Poisson solver (James 1977) is bundled verbatim with my code, and my students Neil Raha, David Earn, Victor Debattista, and Juntai Shen. The SCF option in the code originates from Hernquist & Ostriker (1992) and I am grateful for their permission to include some of their code. Joel Berrier has assisted by testing pre-release versions of the package. This work was also supported in part by NSF grant AST/1108977.

REFERENCES

- Debattista, V. P. & Sellwood, J. A. 2000, *ApJ*, **543**, 704
 Earn, D. J. D. & Sellwood, J. A. 1995, *ApJ*, **451**, 533
 Hernquist, L. & Ostriker, J. P. 1992, *ApJ*, **386**, 375
 James, R. A. 1977, *J. Comp. Phys.*, **25**, 71
 James, R. A. & Sellwood, J. A. 1978, *MNRAS*, **182**, 331
 Swarztrauber, P. N. 1982, in *Parallel Computations*, ed. G. Rodrigue (Academic Press), p51
 Sellwood, J. A. 1981, *A&A*, **99**, 362
 Sellwood, J. A. 1985, *MNRAS*, **217**, 12
 Sellwood, J. A. 2003, *ApJ*, **587**, 638
 Sellwood, J. A. & Athanassoula, E. 1986, *MNRAS*, **221**, 195
 Sellwood, J. A. & Debattista, V. P. 2006, *ApJ*, **639**, 868
 Sellwood, J. A. & Valluri, M. 1997, *MNRAS*, **287**, 124
 Springel, V. 2005, *MNRAS*, **364**, 1105
 Stadel, J. G. 2001, PhD. thesis, University of Washington.